

A Proposal for Secure Access to a NOAA Grid

May 17, 2005

I. Overview

As part of the current Research and Development High Performance Computer System (RDHPCS) procurement, there is a push to provide simple, uniform, secure access to the various HPCS sites (currently located at GFDL, NCEP and ESRL). The hope is that these systems can be linked together into a grid such that a single sign-on provides access to all systems and that jobs submitted at a given site can execute on any of the grid systems. In order to make this possible, the security teams at each site must be able to trust the security measures utilized by the other sites. Ultimately, the government-wide FIPS 201 plan should enable this level of trust. However, implementation of FIPS 201 will be months or much longer away. In this proposal, we outline a near-term approach that we originally hoped would provide a level of trust sufficient to allow the current HPCS's to be connected into a developing NOAA grid relatively quickly so that we can begin to tackle the issues of how to manage it. However, an internal review uncovered a couple of significant obstacles discussed in section IV. For now, we ask the reader to consider how these obstacles can be overcome.

II. Requirements

1. For a node to be added to the grid, it must only be available via secure access technology that **ALL** grid site security administrators trust. As examples, ESRL currently utilizes the Safeword token card technology and GFDL uses Cryptocards. Both of these cards employ one-time passwords and require PIN numbers. Initially, it is envisioned that only the HPCS's at ESRL (Ijet and Ejet) and GFDL (Altix and O3K) would be part of the grid. **In other words, we would be creating a secure perimeter around these systems.**
2. Grid network traffic will only occur between these systems and will be allowed over a defined set of port numbers. We envision these port numbers would be as follows:

a. 2119

The Globus Gatekeeper (see <http://www.globus.org/>)

- b. 2811 The Globus Grid FTP Server
 - c. 8080 Web-services
 - d. 40000-41000 Ephemeral port range
3. Secure access to grid nodes will be enabled by use of the Globus Grid Security Infrastructure (GSI). A description of GSI can be found at <http://www.globus.org/>. Public key cryptography forms the basis for GSI.
 4. A certificate authority (CA) constructed from the Globus simpleCA package will be used to manage grid credentials. A CA administrator will need to be identified. This person will be responsible for verifying if grid certificate requests should be granted and, if so, digitally signing them.
 5. HPCS site administrators will be responsible for providing access to new grid users who have obtained valid CA credentials. In particular, an account would be created and an entry made in the file mapping local user names to grid user "Distinguished Names". Initially, users would be provided access to the individual HPCS sites on an "as-needed" basis. So a particular user may only have access to the GFDL and ESRL machines but not the NCEP machines. However ultimately, we may want to grant universal NOAA grid access as is currently done with the NSF TeraGrid.
 6. The MyProxy online credential repository (<http://grid.ncsa.uiuc.edu/myproxy>) will be used to manage the grid certificates and provide users with short-duration certificate proxies. A MyProxy repository will be located on an extremely securely locked-down node. The only services provided to general users by this node will be the ability to request a grid certificate and the capability to download a certificate proxy. These services will only be available to a user that is logged onto a grid node. A user will only be able to download a certificate proxy. The actual certificates will never leave the repository node. Since the MyProxy repository represents a significant target for hackers, we may want to additionally require that the certificate services be available via token card access. (That is, the user would use a token card to ssh to a grid node and use the token card a second time to access grid proxies from the MyProxy server). Grid certificates will have long lifetimes (1 year?). Proxy certificates will have short durations (12 hours?). To handle longer grid jobs, we may want to allow proxies to be renewed manually or even by authorized grid services on behalf of the user. For example the EU DataGrid project gave its workload management system (WMS) authority to automatically renew credentials of long-running jobs. Ultimately when all grid users have the same secure access technology there should be a single MyProxy repository. However for the interim, we envision a MyProxy repository at each HPCS site.

III. Example Use Case

A user wants to logon to Ijet but run a job remotely on the GFDL Itanium cluster (IC). After using her token card to get access to Ijet she runs a client command that contacts the MyProxy repository service located on the ESRL MyProxy server (call it eproxy). She provides the service with some identifying information (name, phone number, ESRL sponsor, etc.). The CA administrator checks that she is listed in the NOAA locator and then phones her to verify that she has submitted this request. He also verifies with the ESRL sponsor that this user needs access to the NOAA grid. Assuming her request is valid, a grid certificate is created for the user and stored on eproxy. He then notifies her that grid access has been granted; providing her with her grid Distinguished Name. After receiving notification that grid access has been granted, she provides the distinguished name to the ESRL MyProxy server and downloads a certificate proxy. She is now signed on to the grid and can remotely login to IC using a Globus tool called “glogin” and compile her model. She can bring some model data files from Ijet to IC using globus-url-copy. Finally, she can submit jobs to Ijet and IC using a grid scheduler currently under development (for example, a model ensemble where 5 instances run on Ijet and 5 on IC).

IV. Issues

1. At both GFDL and ESRL, various file-systems are mounted via NFS with a variety of servers and personal workstations throughout the respective laboratories. In order to provide a secure token card based perimeter around these HPCS's, the NFS mounts would have to be eliminated. One alternative would be to provide file-systems that securely mirror those located on the HPCS's. However, there are cost and performance issues associated with this approach.
2. From workstations inside the FSL and GFDL firewalls, it is currently possible to gain access to the respective HPCS's without a token card. This access method would have to be eliminated.
3. Login session time limits would have to be instituted. Otherwise, glogin sessions could be inadvertently left unattended; leaving the grid vulnerable for times exceeding the standard certificate proxy durations.

All of these issue beg the following question: **Grid or not, are these good security policies anyway?**